

# Grid Data Management

## Distributed Database Management and Replication with Object Granularity

Koen Holtman, Nov 2, 2000

In the context of the ALDAP (Accessing Large Data Archives in Astrophysics and Particle Physics) project funded by the NSF KDI program, we are collaborating with the SDSS Science Archive team at JHU [1] to exchange knowledge and experience in implementing large science archives using Objectivity/DB. The JHU group has extensive experience in database structures, optimizing database queries, and configuring small and medium size data servers for high throughput, which are applicable to R&D for CMS.

**Apr 2000:** Started working at Caltech.

**Jul 2000:** SDDS SX software installed at Caltech.

The team at JHU is responsible for developing the SX tool and server which are used to analyze SDDS data. Development copies of the SX query tool and server [2] have been successfully installed at Caltech in July 2000. This is a first step towards installing a complete replica of the production SDSS SX data at Caltech.

**Jul 2000:** Finished design of prototype object replication tool.

We aim to develop software tools and middle-ware that allows for large scientific datasets to be managed at the granularity level of single objects, and to be transported across networks as user-defined object collections. The objects (and collections) to be managed can be as small as 100 bytes, with a projected median size around 10 KB [4]. The desire to manage and replicate data at this fine level of granularity is driven by the observation that, in LHC physics datasets, the fundamental unit of information is an event, an observation record which has in the order of 1 MB of data associated with individual objects (grouped in containers). A design for a first prototype was completed in July 2000 [5] [6]. The prototype will use Globus components and tools such as Globus GSI FTP [3], and will integrate with ORCA production data. In August 2000 a decision was made to re-use parts of the GDMP software in developing the prototype.

**Sep 2000:** Installation of Globus, GDMP, ORCA\_4\_3\_0\_pre4, IGUANA, preparing for use of these in software development.

**Oct 2000:** Finished implementation of core components of prototype object replication tool.

A talk about this was given at ACAT2000 [7].

**Nov 2000:** Demo of prototype object replication tool at SC2000 using Globus middleware ORCA data.

The demo is titled 'Bandwidth Thirsty Particle Physics Event Collection Analysis and Visualization Using Object Databases and the Globus Grid Middleware'. We will demonstrate high bandwidth retrieval of particle physics event collections from an object database hosted at Caltech. The application we will use is based on Globus middleware and the Objectivity ODBMS, and will rely heavily on optimized network paths between Dallas and both Caltech and CERN. The main goal of the demo is to show (to fellow collaborators in the various Grid projects) that object level replication (transporting extracted object sub-collections) is feasible, and to establish its place in the grid architectures being developed by projects like Grifyn.

**Feb 2001:** Creation of complete replica of SDDS analysis dataset at Caltech.

In the context of the ALDAP project we are collaborating with the SDSS Science Archive team at JHU [1] to exchange knowledge and experience in implementing large science archives using Objectivity/DB. The replica of the SDDS analysis dataset, will be used for performance and scalability studies with respect to advanced data analysis and reclustering techniques, and these aspects will be compared to the CMS HLT production federations. From Feb 2000 onwards the Tier2 RC hardware cluster at Caltech will be used with these datasets to drive the development of techniques that exploit the latent massive I/O capacity of the local disks on such clusters. The use of local disks for various tasks will be compared to the use of high-performance centralized RAID arrays.

**Mar 2001:** Proof-of-concept prototype of application of object reclustering and replication techniques in CMS/HLT and SDSS data analysis.

This is planned collaborative work in the ALDAP project, collaborating with either or both of the teams at FNAL and JHU.

**Apr 2001:** First deployment of object level replication services in tier2 prototype centre at Caltech.

This milestone will involve integration of object replication tool with ORCA, Tags, the Globus Replica Catalog, the Globus Replication Manager (if available by that time), and GDMP.

**May 2001:** Proof-of-concept prototype of the use of advanced tag searching techniques developed by Kurt Stockinger, together with navigation back to AOD and fully reconstructed data for some events.

This builds on the work of Kurt Stockinger [8] [9], taking it into deployment on a large computing installation. Both CMS and SDDS data will be used for the proof-of-concept study. This is joint work with Kurt Stockinger of CERN/IT, who will visit Caltech starting in March.

**September 2001:** Complete design for the integration between site-internal task scheduling and site-internal data replication.

We expect this activity to have a strong interaction with the various Grid projects, in particular with the research and design activities in Grifwyn.

**Dec 2001:** Integration between task scheduling and data replication, producing a pilot installation that works in terms of virtual data.

## References

- [1] <http://www.sdss.jhu.edu/>
- [2] <http://www.sdss.jhu.edu/ScienceArchive/doc.html>
- [3] <http://www.globus.org/datagrid/deliverables/gsiftp-tools.html>
- [4] K. Holtman, H. Stockinger. Building a Large Location Table to Find Replicas of Physics Objects. Proceedings of CHEP 2000, Padova, Italy.  
[http://kholtman.home.cern.ch/kholtman/olt\\_long.ps](http://kholtman.home.cern.ch/kholtman/olt_long.ps)
- [5] [http://www.cacr.caltech.edu/ppdg/meetings/ppdg\\_collab/holtman/objrepl.pdf](http://www.cacr.caltech.edu/ppdg/meetings/ppdg_collab/holtman/objrepl.pdf)
- [6] [http://kholtman.home.cern.ch/kholtman/globusretreat\\_objrepl.ppt](http://kholtman.home.cern.ch/kholtman/globusretreat_objrepl.ppt)
- [7] K. Holtman. Object level physics data replication in the Grid ACAT'2000, October 16-20, 2000, Fermi National Accelerator Laboratory  
<http://conferences.fnal.gov/acat2000/program/presentations/vlsc302.ppt>
- [8] <http://kurts.home.cern.ch/kurts/PHD/Phd1.htm>
- [9] Kurt Stockinger, Dirk Duellmann, Wolfgang Hoschek, Erich Schikuta. Improving the Performance of High Energy Physics Analysis through Bitmap Indices. 11th International Conference on Database and Expert Systems Applications, London - Greenwich, UK, Sept. 2000. Springer-Verlag.  
[http://kurts.home.cern.ch/kurts/RESEARCH/Bmi\\_dexa2000.ps](http://kurts.home.cern.ch/kurts/RESEARCH/Bmi_dexa2000.ps)